

(43) Date of A publication 22.12.1993

(21) Application No 9312417.0

(22) Date of filing 16.06.1993

(30) Priority data

(31) 901454

(32) 19.06.1992

(33) US

(71) Applicant

Westinghouse Electric Corporation

(Incorporated In the USA - Pennsylvania)

Westinghouse Building, Gateway Center, Pittsburgh,
Pennsylvania 15222, United States of America

(72) Inventors

Warren Albert Edblad

Albert William Crew

Carl Joseph Staab

Gilbert William Remley

Linda Lee Santoline

(74) Agent and/or Address for Service

Ronald Van Berlyn

23 Centre Heights, London, NW3 6JG,
United Kingdom

(51) INT CL⁵

H04L 12/18, G06F 12/08

(52) UK CL (Edition L)

H4P PPG

U1S S1877 S1901 S2047 S2145 S2150 S2184
S2207

(56) Documents cited

GB 2205018 A EP 0303830 A2

(58) Field of search


UK CL (Edition L) H4P PPG PPS

INT CL⁵ H04L 12/18

(54) Transmission and storage of real time data in a network system

(57) A computer network system (1) repetitively distributes uniquely identified blocks of real time data over a communications network (5) to all real time stations (3a, 3b) for storage of each repetition of each entire block of data directly in station memory (15) at a unique address space assigned to that uniquely identified block of data.

Application is to e.g. a process control system with input sensors such as thermocouples (7) and output contacts (9) for a valve or motor. Man-machine interface stations (3b) include a workstation (19) with CRT display (21). Other applications include automatic teller machines, video and multi-media mixing networks.



11/8




FIG. 1




FIG. 2




FIG. 3

3/8.




FIG. 4A

4/8




FIG. 4B

5/8

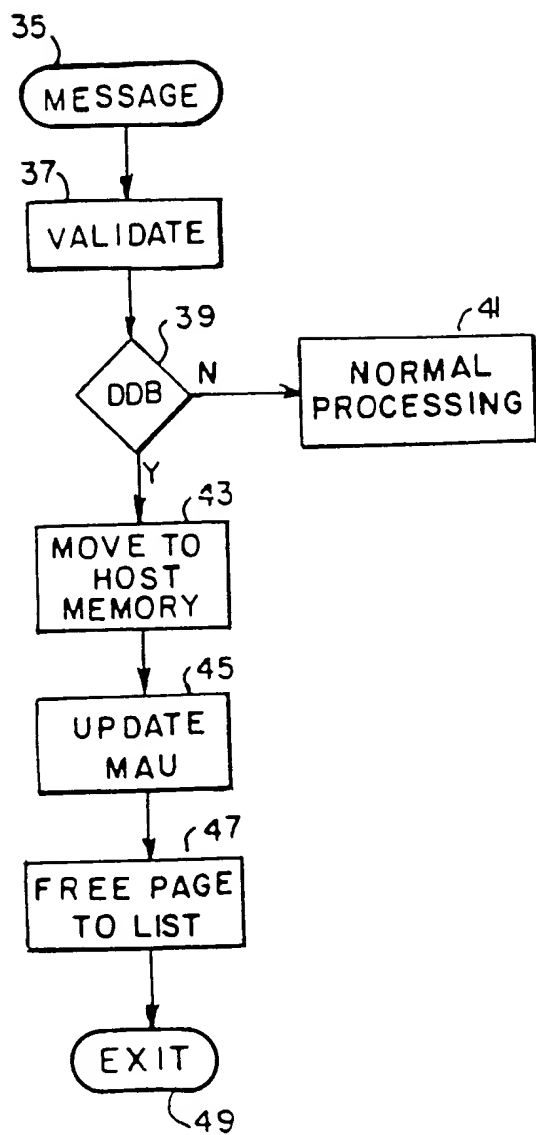


FIG. 5

6/8




FIG. 6A




FIG.6B

8/8

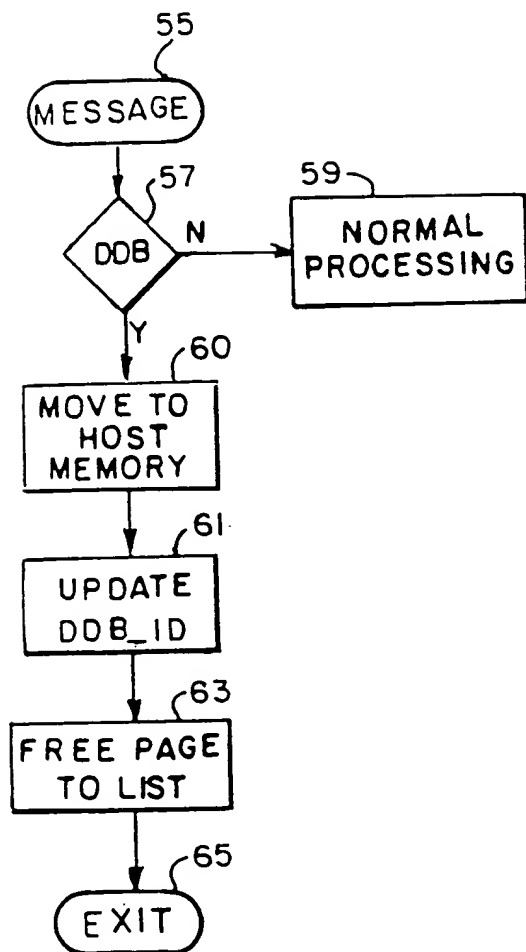


FIG. 7




FIG. 8

A REAL TIME DATA IMAGING NETWORK SYSTEM
AND A METHOD OF OPERATING SAME

This invention relates to sharing of data in a distributed processing system, and more particularly to the transmission, processing and storage of real time data in a network system.

5 In most computer networks, data transmitted over a data highway interconnecting the various stations in the network is addressed to a specific station or stations. The message is ignored by the remaining stations on the network. The network interface at the addressed station
10 interrupts the station central processing unit which then processes the message. Often the receiving station sends back an acknowledgment to the sender station. In these systems, the message identifies both the receiving and sending stations. Such a process is efficient in an
15 application with interactive communication needs between stations.

 In broadcast multi-cast networks, the messages are not addressed to any specific station. Each station
20 receives all of the data transmitted and traditionally takes what it needs. This also requires processing of the message by the central processing unit of the receiving station. Typically, to reduce the burden on the system, data is only broadcast by exception. That is, only when the value of a piece of data changes is a message sent.
25 Thus, for instance, in a distributed process control system, where various functions are being performed by a

number of stations on a network, data is only broadcast when the value of a variable changes. Other deterministic systems, such as automatic teller machines linked to a central unit by a network also only send data when there is a transaction to process. However, such systems can become overburdened if there is a large disturbance in the process being controlled, or a large number of transactions occur simultaneously.

Assignee of the present invention has developed a broadcast network in which data from the various stations are bundled into packets or blocks of data which are transmitted periodically, whether or not the data has changed. Data with a short latency is transmitted at a fast repetition rate, while other data which is less critical or changes at a slower rate, is transmitted at more widely spaced time intervals. In this system, all the data is transmitted repetitively and is received by all of the stations. Each station examines each block of data received and slices off for storage only that data which it needs. This system is designed for the worst case scenario, and thus does not become overwhelmed by large disturbances in the system. However, local processing of each block of data is required to determine the data to be stored in memory at that station, which limits the data rate that can be handled by the system.

There is a need, therefore, for an improved broadcast network system with an enhanced data rate and yet which does not become overwhelmed by the worst case scenario. Additionally, peer-to-peer communications should also be provided without compromising the real-time performance.

While other systems attempt to improve performance by limiting the data transmitted by the sending station, or limiting the data stored at the receiving station, the present invention takes the novel approach of transmitting all of the real time data to all of the stations requiring real time data, and storing all of the data transmitted at each such station, whether all of the data is needed or not at a specific station. While at first glance this may

appear to be inefficient, in reality, through put is improved because minimal processing is required at each end of the transmission. At the sending end, the entire block of data is transmitted. No processing is required to
5 determine what values have changed. At the receiving end no processing is required to determine what data to store.

Rather than identifying each block of data by the sender or receiver station, each block is given a unique identifier. Each receiving station stores subsequent
10 repetitions of each uniquely identified block of data in the same uniquely identified address space at the station. The unique address space is preferably a virtual address space. Also, preferably, virtual address spaces are translated into physical address spaces at the local station by
15 memory management hardware which maintains a list of free pages in memory and inserts the incoming message image in the next available page and then updates the translation of the virtual address space to this new physical location. With this arrangement, the local central processing unit
20 can by accessing the last image of a message while the current message is being stored. This process places minimum burden on the host central processing unit.

Alternatively, in systems without appropriate memory management hardware, a driver routine can be used to
25 maintain a table for virtual to physical address space translation and a separate list for free pages in memory. However, even this alternative places a very small processing load on the local central processing system.

The invention embraces both a system for and
30 method of implementing data transfer in a real time data imaging network.

The object of the invention is to provide an efficient data highway for applications that perform broadcast functions.

35 With this object in view, the present invention resides in a real time data imaging network system comprising: a plurality of real time stations, each having memory means; characterized by a broadcast communications network

connecting said stations and repetitively transmitting from at least one of said stations to each other station at least one uniquely identified block of data with each repetition containing a current data image, said station
5 including storing means storing each repetition of said uniquely identified block of data in a same unique address space in said memory means.

A full understanding of the invention can be gained from the following description of the preferred
10 embodiments when read in conjunction with the accompanying drawings in which:

Figure 1 illustrates schematically a real time data imaging systems in accordance with the invention.

Figure 2 illustrates the definition of uniquely
15 identified blocks of data which are transmitted throughout the system of Figure 1 in accordance with the invention.

Figure 3 schematically illustrates the operation of a memory addressing unit, which forms part of the system of Figure 1, in translating a virtual memory address space
20 into a physical memory address space.

Figures 4A and 4B illustrate storing in a receiving station memory two successive images of a specific block of data transmitted by a network which forms part of the system.

Figure 5 is a flow chart for a computer program
25 for a network interface at a receiving station which implements the process of Figures 4A and 4B.

Figures 6A and 6B illustrate another embodiment of the invention for storing in a receiving station two
30 successive images of a uniquely identified data block transmitted over the network.

Figure 7 is a flow chart for a suitable computer program for the station central processing unit which implements the process of Figures 6A and 6B.

Figure 8 is a flow chart for a suitable computer
35 program for the station central processing unit for accessing data stored by the embodiment of Figures 6A, 6B and 7.

The invention is directed to a real time data imaging network system. It has particular application, and will be described as applied to, a distributed process control system. However, it will be appreciated by those skilled in the art that the invention has application to other deterministic network systems with a distributed data base, such as for instance, automated teller machines, video and multi-media mixing networks and others.

Figure 1 illustrates a distributed process control system based on the real time imaging network 1. The system 1 comprises a number of stations 3 interconnected by communication network 5. The stations, or "drops", as they are often referred to, include, but are not limited to, data processing units (DPUs) 3a man-machine interface units (MMI) 36. Other drops (not shown) that do not require real time data can be included on the networks. These DPUs 3a are typically connected to physical devices associated with the process to be controlled. The physical devices include: input sensors, such as thermocouples 7, and outputs, such as contacts 9. Each DPU 3a provides control over loops assigned to it by monitoring the operation of each loop and insuring that the values of the pertinent variables are kept to their respective set points. The DPU does this by calling algorithms which scan the input signals from the sensors, such as 7, compare the variable actual value to its set point, perform the calculations necessary to bring the variable values into line with the set point, and output a control signal to a final control device at the process such as the contacts 9 which for instance, may control a valve or a motor, etc. Thus, items in a physical environment such as a coolant level or fuel flow are understood by the control system as data. Information about a process variable is stored in the DPU 3a as a point data record. The information in this data record can be broadcast onto the communications network 5 for other stations or drops to use, or it can be kept internal to the drop to which it is assigned. A point data record contains two distinct types of information. One type is

the dynamic information, which is a function of the physical inputs of the drop. An example of dynamic information is the process flow value read from a flow sensor 7. The other type of information is the static information, which includes configuration information and other data which rarely changes. Examples of static information are the conversion constants, alarm limits, and textural descriptions associated with an analog input. While the dynamic information must be broadcast periodically, the static data is broadcast on an "as needed" basis.

The DPUs 3a include a central processing unit (CPU) 11 which receives the input information from the sensor 7, implements the control algorithms, and controls the output devices 9. The DPUs 3a also include a network interface 13 between the CPU 11 and the network 5. The DPUs 3a further include a memory 15 which is shared between the CPU 11 and the network interface 13.

A memory address unit (MAU) 17 provides virtual to physical address mapping for the memory 15 in which data is stored and retrieved. The shared memory 15 and MAU 17 may be physically provided in the CPU 11 or the network interface 13.

The man-machine interface (MMI) stations, or drops, 3b, include a workstation 19 which incorporates a CRT 21 for display of information and a CPU 23. The drops 3b also contain a network interface 13 and shared memory 15 with its MAU 17 which again, can be physically provided in the workstation 19 or the network interface 13.

The communication network 5 includes coaxial, twisted pair, or fiber optic, cables 25 and 27 strung throughout the facility between the stations and configured to provide redundant communication paths. In Figure 2, this is illustrated as a dual counter-rotating ring.

Each of the drops 3 must be autonomous, that is, it must be capable of functioning without the presence of the communications network 5. This imposes a requirement that the database for the system be distributed. For a DPU 3a, all its configuration data must be stored internally at

the drop, so that the DPU can restart and perform its control function without the need for data from any external source. For an MMI 3b, the configuration data must also be stored internally so that the drop can restart.

5 However, since the primary function of an MMI is to serve as a man-machine interface, the workstation requires the presence of the communications network 5 to obtain the process data associated with its display screens.

The preferred communications network highway 5 has

10 an open system architecture conforming to the FDDI (fiber distribution data interchange) standards. The communications network 5 provides both proprietary and general purpose communications. General purpose communication is managed via a recognized standard protocol, such as the

15 TCP/IP or UDP/IP protocol of the Internet Protocol Suite, encapsulated within an FDDI asynchronous message. The general purpose communications protocol is used whenever the performance is adequate. In contrast, the proprietary communications protocol to be described is used to handle

20 message traffic where the data latency requirements prohibit the use of the higher level protocol, such as for dynamic point data communications.

System data elements that are transmitted over the communications network 5 are referred to as network points.

25 Data elements which are local to a drop are referred to as local points. Dynamic data associated with periodic, network points are transmitted periodically at either a SLOW RATE or a FAST RATE. Due to the large number of dynamic data points in a system and the overall through put

30 requirements of the dynamic point data, the software overhead associated with message processing prevents the transmission of dynamic point data associated with periodic, network points as individual messages. For this reason, the dynamic data associated with periodic, network

35 points is bundled into dynamic data blocks (DDBs) for transmission on the highway. A dynamic data block is associated with the drop which originates the points contained within the block, and is referenced by a numeri-

cal identifier which is unique throughout the system. Each dynamic data block has a fixed transmission rate either (slow or fast) and a fixed size shared memory area associated with it where it is stored. As periodic, network points are added to each drop 3, they are assigned to a dynamic data block of the correct transmission rate which is originated by that drop. Bounding of the maximum dynamic data latency is critical to the correct operation of the system. For this reason, transmission of dynamic data blocks is handled as synchronous FDDI frames, and is handled by the network interface 13 without intervention from application software. All general purpose communications are handled as asynchronous FDDI frames.

The definition of a dynamic data block (DDB) 29, is illustrated in Figure 2. The DDB 29 consists of a field containing a unique identifying index (DDB_ID), a fixed size header portion used for control and status, and a data portion containing the assigned dynamic data points (DDPs). The header includes an age counter which is a measure of the age of the DDB. The originating station sets the age counter to an age threshold for the DDB. An application process running in each receiving drop periodically decrements this field. If it reaches zero, the block is considered "stale" and will not be used.

The active length portion of the header contains the number of bytes of the DDB containing DDPs, plus the size of the header and ID fields. The version portion of the header is updated whenever DDPs are added, modified, or deleted from the DDB.

A unique feature of the invention is that the dynamic data blocks, DDBs, generated by each real time station are broadcast over the communications network to, and entered in storage in their entirety in, each other real time station. While other broadcast systems have attempted to increase efficiency by reducing through put, this invention actually transmits all data to all real time stations whether they need all the data points or not. The key factor is that this approach requires no, or minimal,

processing of the communication by the CPU of the receiving station. This is accomplished by directly storing repetitions of each uniquely identified block of data in a uniquely identified address space in each real time station so that the uniquely identified address space contains the current, or most recent, image of that block of data. The unique address space is assigned by the local station and does not necessarily have to be the same for each station.

When the uniquely identified dynamic data block is received at a station by the network interface 13, it is distinguished by the network interface from the lower priority messages. The network interface which contains a processor, maintains the list of unique addresses for that station for each of the uniquely identified dynamic data blocks. In a preferred form of the invention, these unique addresses for the uniquely identified dynamic data blocks are virtual addresses, and not absolute physical addresses in the memory 15.

The preferred implementation of the invention minimizes the impact on the host processor by utilizing the memory manager and hardware provided in most RISC (reduced instruction set computer) workstations today. This implementation assumes that the host processor 11, 23 uses the hardware memory address unit 17 to implement the translation of virtual memory address spaces into the corresponding physical memory address spaces. In general, the MAU 17 can be thought of as an array of hardware registers 31 which are indexed by a portion of the virtual memory address. As shown in Figure 3, the indexed register provides the page address (A_n to A_{12} in the example of Figure 3) of the physical memory which the hardware combines with the remaining part of the virtual address (A_{-11} - A_0 in the example) to form the physical memory address. Typically MAUs 17 operate with a page size of 4096 bytes which matches nicely with a DDB size of 4K, wherein each DDB occupies a full page. The DDB size can be adjusted to match the page size or multiple pages can be used for each DDB.

Figures 4A and 4B illustrates apparatus 32 and a process for implementing this aspect of the invention through an example in which four DDBs 1-4 are shown stored in pages 33 of the memory 15. For the state shown in Figure 4A, the MAU 17 translates the virtual memory address for the latest image of the block of data DDB 1 to page 45 in the physical memory. The latest images of DDBs 2-4 are stored on pages 43, 48 and 50 respectively in physical memory. The network interface 13 maintains a list of available pages. In the example, those available pages are 47, 49, 44 and 46 in that order. When the next repetition of the first data block, DDB 1a, is received by the network interface 13 automatically stores that current image of the first DDB directly in page 47 of the memory, the first page on the list of available pages.

The network interface 13 then remaps the MAU 17 as shown in Figure 4B to indicate that the most current image of the first data block, DDB 1a, is stored in page 47. As shown, page 45 is added to the end of the available buffer list and 49 moves to the top of the list. The next message image, DDBn, is then stored in page 49 as indicated in Figure 4B. Thus the incoming data is stored in its entirety in the next available page in the memory 15. Minimum processing load is imposed on the network interface 13. No load is placed on the host processor except in instances where updating of the MAU registers 31 may require minimal action by the host processor.

Figure 5 illustrates a flow chart of a suitable computer program run by the network interface 13 to implement the procedure illustrated in Figures 4A and 4B. As indicated at 35 in Figure 5, the network interface 13 is activated by receipt of an FDDI frame. A check is made at 37 to verify the validity of the message. If the message is not a DDB as determined at 39, then the message is passed through at 41 to the standard FDDI protocol suite with no special handling.

If the message is a DDB as determined at 39, the entire frame is transferred into the first memory page at the head of the free buffer list, and that buffer is removed from the free list at 43. Using the DDB unique
 5 identification number from the received message, the corresponding MAU register 31 is updated at 45 to the newly received buffer.

The memory page which was removed from the MAU register is linked to the tail of the free page list at 47,
 10 and the network interface returns control to the host processor while waiting for the next FDDI frame at 49.

An alternative implementation 52 of the invention does not utilize the MAU 17; however, it does impact the performance of the host processor, 11, 23. In this imple-
 15 mentation 52, which is illustrated by example in Figures 6A and 6B, the reception of a new image of a DDB is controlled by the FDDI device driver routine resident in the host processor software. This FDDI device driver is required to maintain the list of free pages in a table 51 from which it
 20 selects the next available page. All application access to the data is made indirectly by obtaining the page base addressed from a table 53 which is indexed by the DDB number. The application process must then compute the appropriate shared memory address by combining the base
 25 address with the offset into the page. This computation must be performed for every read or write to the DDBs.

As shown in Figure 6A. the network interface 13 receiving a new image, DDB 1a, of the first message stores it directly in shared memory at base address 4000_{16} , the
 30 first free page on the free page list 51. During this operation, an application process is accessing the data contained in the last image of the first message, DDB1. This access is accomplished by indexing the DDB_ID table 53 to obtain the current base address (2000_{16}) computing the
 35 offset from this, and accessing a shared memory page 33 for the data. When the reception of DDB 1a is complete, the device driver must substitute the base address of the newly received DDB 1a image with the page in the DDB_ID table 53

index 1 as shown in Figure 6B. Page 2000₁₆, the address of the previous image, DDB 1, of the first message, is added to the bottom of the free page list 51. Figure 6B also shows another message image, DDBn, being stored by the network interface 13 in page 6000₁₆, the address which moved to the top of the free page list 51.

Figure 7 illustrates a flow chart of the host CPU in carrying out the above steps of the alternate embodiment of the invention. The driver software in the host CPU is activated by the receipt of a valid FDDI frame as indicated at 55. The validation is performed by the network interface 13. If the multi-cast address and the unique protocol identifier indicate at 57 that this is not a DDB packet, the non-DDB frame is passed through to the standard protocol suite with no special handling as indicated at 59.

If this is a DDB, as determined at 57, the entire frame is transferred into the memory page at the head of the free buffer list, and that buffer is removed from free list at 60. The contents of the DDB_ID table 53, at the location indexed by the DDB_ID field from the received frame, is changed to the address of the page where the frame was just stored as indicated at 61. The memory page which was replaced in the DDB_ID table 53 is linked to the tail of the free page list as indicated at 63. The routine then exits at 65 to wait for the next FDDI frame to reactivate the driver.

For the embodiment of the invention illustrated in Figures 6A-B, 7, the application software uses the routine shown in flow chart form in Figure 8 to access the DDB contents. Hence, when the application program calls for access to DDB contents as indicated at 67, the DDB's unique ID is used to index the DDB_ID table 53 as indicated at 69. Using the address obtained from the DDB_ID table 51, the offset within that data block for the particular parameter value required by the application process is added to obtain the actual memory address as indicated at 71. The computer memory address is then accessed at 73 to obtain

the required data which is returned to the application process at 75.

5 While specific embodiments of the invention have been described in detail, it will be appreciated by those skilled in the art that various modifications and alternatives to those details could be developed in light of the overall teachings of the disclosure. Accordingly, the particular arrangements disclosed are meant to be illustrative only and not limiting as to the scope of the invention
10 which is to be given the full breadth of the appended claims and any and all equivalents thereof.

IDENTIFICATION OF REFERENCE NUMERALS USED IN THE DRAWINGS

<u>LEGEND</u>	<u>REF. NO.</u>	<u>FIGURE</u>
CPU	11	1
NETWORK INTERFACE	13	1
NETWORK INTERFACE	13	4A
NETWORK INTERFACE	13	4B
NETWORK INTERFACE	13	6A
NETWORK INTERFACE	13	6B
MEM	15	1
MAU	17	1
MEMORY ADDRESSING UNIT	17	3
CPU	23	1
MESSAGE	35	5
VALIDATE	37	5
DDB	39	5
NORMAL PROCESSING	41	5
MOVE TO HOST MEMORY	43	5
UPDATE MAU	45	5
FREE PAGE TO LIST	47	5
EXIT	49	5
MESSAGE	55	7
DDB	57	7
NORMAL PROCESSING	59	7
MOVE TO HOST MEMORY	60	7
UPDATE DDB_ID	61	7
FREE PAGE TO LIST	63	7
EXIT	65	7
ACCESS DDB DATA	67	8

IDENTIFICATION OF REFERENCE NUMERALS USED IN THE DRAWINGS

<u>LEGEND</u>	<u>REF. NO.</u>	<u>FIGURE</u>
INDEX DDB_ID	69	8
GET ADDRESS	71	8
ACCESS DATA	73	8
DONE	75	8

CLAIMS:

1. A real time data imaging network system (1) comprising:

a plurality of real time stations, each having memory means (15); characterized by

5 a broadcast communications network (5) connecting said stations (3) and repetitively transmitting from at least one of said stations to each other station at least one uniquely identified block of data (DDB) with each repetition containing a current data image, said station
10 including storing means (32, 52) storing each repetition of said uniquely identified block of data in a same unique address space in said memory means.

2. The network system of claim 1 wherein a plurality of said real time stations (3) each generate at
15 least one uniquely identified block of data (DDB) and wherein said communications network (1) repetitively transmits each of said uniquely identified blocks of data from each of said plurality of stations to each other station, and wherein said storing means (32, 52) in each station
20 stores each repetition of each uniquely identified block of data in a same address space in said memory means (15) unique to that uniquely identified block of data.

3. The network of claim 2 wherein at least some of said stations (3) generate first and second sets of
25 uniquely identified blocks of data (DDB), and wherein said communications network (1) repetitively transmits said first set of uniquely identified blocks of data at a first repetition rate and transmits said second set of uniquely

identified blocks of data at a second slower repetition rate.

4. The network system of claim 2 wherein said same address spaces are unique virtual address spaces and wherein said storing means (32, 52) includes means (17, 53) translating said unique virtual address space for each uniquely identified block of data into a physical address space in said memory means in which a current repetition of said uniquely identified block of data is stored.

5. The network system of claim 2 wherein each station (3) includes a central processing unit (11, 23), wherein said unique address spaces are unique virtual address spaces and wherein said storing means (13, 51) include means maintaining a list of available physical address spaces in said memory means, means storing a uniquely identified block of data received over said communications network in a selected one of said available physical address spaces in said memory means, and means (17, 53) translating said virtual address space into the selected physical address space for access to said block of data by said central processing unit (11, 23).

6. The network system of claim 5 wherein said means maintaining said list of available physical address spaces and said means storing a uniquely identified block of data in a selected physical address space are provided in a network interface (13) and wherein said translating means is a memory address unit (17).

7. The network system of claim 5 wherein said means maintaining said list of available physical address spaces in said memory means and said means translating said virtual address space into said selected physical address space are provided in said central processing unit (11, 23) and wherein said means storing said uniquely identified block of data received over said communications network (5) is a network interface (13).

8. A method of distributing data in a real time data imaging network system (1) comprising; a plurality of real time stations (3) each containing a central processing

unit (11, 23), a memory (15) and a network interface (13); and a communications network (5) interconnecting said stations (3) through said network interfaces (13), said method comprising:

5 operating said central processing unit (11, 23) in each station (3) to repetitively assemble data into uniquely identified blocks of data (DDB) containing a most recent data image;

 broadcasting each repetition of each uniquely
10 identified block of data (DDB) assembled by each station over said communications network (5) to every other station; and

 storing in said memory (15) at each station each
15 repetition of each uniquely identified block of data received from other stations over said communications network in a same unique address spaced in said memory (15) for each uniquely identified block of data.

 9. The method of claim 8 wherein said unique
20 address spaces are unique virtual address spaces and wherein said storing step comprises:

 storing each uniquely identified block of data
(DDB) received from other stations (3) over said communications network (5) in an available physical address space
(33) in said memory and translating said available physical
25 address space into said unique virtual address space for that uniquely identified block of data.

 10. The method of claim 9 including following
translating the selected available address space into the
unique virtual address space of the uniquely identified
30 block of data (DDB) stored at that physical address space, adding a physical address space at which a previous repetition of that uniquely identified block of data was stored to a list (51) of available physical address spaces.

 11. The method of claim 10 wherein said list (51)
35 of available physical address spaces is maintained in said network interface (13), wherein said network interface stores said uniquely identified blocks of data (DDB) received over said communications network (5) in said

memory (15) and wherein translating of said virtual address spaces into said selected physical address is accomplished by a memory address unit (17).

5 12. The method of claim 10 wherein said central processing unit (11, 23) at each station is programmed to maintain said list (51) of available physical address spaces and to translate said virtual address space into said selected physical address space, and wherein said
10 storing of said uniquely identified block of data received over said communications network (5) in a selected physical address space is performed by said network interface means (13).

Examiner's report to the Comptroller under
Section 17 (The Search Report)

Application number

GB 9312417.0

Relevant Technical fields

(i) UK Cl (Edition L) H4P (PPG, PPS)

(ii) Int Cl (Edition 5) H04L 12/18

Search Examiner

K WILLIAMS

Databases (see over)

(i) UK Patent Office

(ii)

Date of Search

31 AUGUST 1993

Documents considered relevant following a search in respect of claims 1-12

Category (see over)	Identity of document and relevant passages	Relevant to claim(s)
A	GB 2205018 A (REUTERS) - see abstract	1, 8
A	EP 0303830 A2 (IBM) - see column 10, lines 20-51 and US 4807224	1, 8

Category	Identity of document and relevant passages	Relevant to claim(s)

Categories of documents

X: Document indicating lack of novelty or of inventive step.

Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.

A: Document indicating technological background and/or state of the art.

P: Document published on or after the declared priority date but before the filing date of the present application.

E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.

&: Member of the same patent family, corresponding document.

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).